

Technical Report No. 5
on
ANALYSIS AND DEVELOPMENT
OF
FAILURE-RESPONSIVE SYSTEM ORGANIZATIONS

Contract NASw-572
Reference WGD-38521

December 1964

FACILITY FORM 602

N65 17605	
(ACCESSION NUMBER)	(THRU)
76	1
(PAGES)	(CODE)
CR 60897	10
(NASA CR OR TMX OR AD NUMBER)	(CATEGORY)

GPO PRICE \$ _____

OTS PRICE(S) \$ _____

Hard copy (HC) 3.00

Microfiche (MF) .75



Westinghouse Defense and Space Center
Surface Division

P. O. BOX 1897

Baltimore, Md., 21203

Technical Report No. 5
On
Analysis and Development of Failure-Responsive
System Organizations

Contract NASw-572
Reference WGD-38521

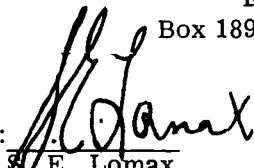
by
C. G. Masters, Jr.

December 1964

Note: This report is presented in the form of a thesis. As a thesis, the report was submitted by the author to the University of Pittsburgh in partial fulfillment of the requirements for the degree of Master of Science in Electrical Engineering in December 1964.

The Westinghouse Electric Corporation
Electronics Division
Box 1897, Baltimore 3, Maryland

APPROVED:


E. Lomax,
Advanced Development Engrg.

MDE 4472

ABSTRACT

THE ANALYSIS AND DEVELOPMENT OF FAILURE RESPONSIVE REDUNDANT SYSTEM ORGANIZATIONS

17605

The problem of achieving high reliability in electronic systems has become increasingly difficult to solve as the systems have decreased in physical size and increased in functional complexity. The importance of solving this problem has also grown with the evolution of new applications which often make these systems vital to national security and human welfare.

Several methods for employing redundant equipment to achieve high system reliability have been proposed. Most of these methods restrict the operation of redundant components or subsystems to a single location within the overall system. The purpose of this thesis is to present a new technique which combats the effects of component failure patterns by allowing redundant subsystems to be shifted around within a system in response to the existing failure pattern. This technique permits the user to more effectively employ redundant equipment in his efforts to increase the useful life of electronic systems.

Author

ACKNOWLEDGEMENT

The development of the concept of failure responsive redundant systems was initially undertaken with the encouragement and support of the Surface Division of the Westinghouse Electric Corporation. The work described in this paper was done under Contract NASw-572 with the National Aeronautics and Space Administration. I should like to acknowledge the generosity and confidence of these organizations in providing support for this work.

I should also like to thank Mr. W. C. Mann for his invaluable suggestions and encouragement; Mr. J. M. Hannigan for his aid in preparing a complex simulation program and in compiling the results obtained from the use of this program; Mr. H. Retsky for his extensive editorial assistance; and Dr. T. W. Sze of the University of Pittsburgh for his aid in the preparation of this thesis.

In addition, I should like to acknowledge the comments and suggestions of the members of the oral examination committee. This committee consists of Dr. T. W. Sze, Chairman, Dr. R. P. O'Shea and Dr. Z. H. Meiksin all of whom are members of the Electrical Engineering Department of the University of Pittsburgh.

TABLE OF CONTENTS

	Page
ACKNOWLEDGEMENT	ii
I. INTRODUCTION	1
A. The Need for High System Reliability	1
B. Methods of Increasing System Reliability	2
1. Conservative Design	2
2. Hyper-reliable Components	2
3. Coding	3
4. Redundant Equipment	3
C. Redundancy Techniques	4
II. THE PURPOSE OF THIS STUDY	9
III. PREVIOUS WORK IN THIS AREA BY OTHER INVESTIGATORS	10
IV. FAILURE RESPONSIVE SYSTEM ORGANIZATIONS	12
A. The General Concept	12
B. The Specific Organizational Objective	15
V. ANALYSIS METHODS.	17
A. The "Brute Force" Method	17
B. The Markov Chain Method	18
C. The Minimal Cuts Method	19
D. The Computer Simulation Method.	21
VI. THE COMPUTER SIMULATION PROGRAM	22
A. The Operational Principles of the Program	22
B. Individual Subsystem Information.	22
1. Failure Location Intervals	22
2. Spare Lists	24
3. Other Stored Data	24

TABLE OF CONTENTS (Continued)

	Page
C. The Detailed Operation of the Program	26
1. Data Storage	26
2. The Simulation Procedure	26
VII. SYSTEM EVALUATION	31
A. Methods for Estimating System Reliability Versus Time Curves. . .	31
1. The Conditional Probability Method.	31
2. The Random Time Generation Method	34
3. Comparison of the Two Estimation Techniques.	37
B. Single-Values Measures of Performance	37
1. Mean Time Between Failures	38
2. System Reliability at a Selected Time.	38
3. Quantile Occurrence	40
VIII. SIMULATION RESULTS	42
A. Phase I Simulations	43
1. Order-Three Systems.	43
a. Experiment I	43
b. Experiment II	45
c. Experiment III	47
d. Experiment IV	48
e. Experiment V	50
2. Order-Four Systems (Experiment VI).	52
3. Functional Order Systems.	54
a. Experiment VII	54
b. Experiment VIII	57
B. Phase II Simulations	58

TABLE OF CONTENTS (Continued)

	Page
IX. SUMMARY AND CONCLUSIONS	63
A. Summary	63
B. Conclusions	63
APPENDIX I	68

LIST OF ILLUSTRATIONS

<u>Figure</u>	<u>Title</u>	<u>Page</u>
1.	Redundant Component Configurations	5
2.	A Segment of an Example System	7
3.	Example Failure Pattern	13
4.	Critical and Non-Critical Order of Failures	20
5.	Two Response Strategies	25
6.	A Typical System and Its Matrix Representation	27
7.	Summary Flow Chart of Computer Program	30
8.	Histogram of Observed System Failures	31
9.	Cumulative History of Observed System Failures	33
10.	Uniform to G(y) Distribution Transformation	35
11.	Comparison of Reliability Estimation Curves	36
12.	Non-Redundant System Reliability Curves	39
13.	Different Reliability Curves with Similar Means.	39
14.	Different System Reliability Curves with Similar Short Life Reliabilities	40
15.	The "Useful Life" Measure	41
16.	Sample Strategies for Consecutive Lists	44
17.	Comparison of Alternating and Sequential Consecutive Lists	45
18.	Comparison of Response Strategies with and without "Rescan" Capability.	46
19.	Sample Strategy for a Normal Set List	46
20.	Comparison of Normal for Consecutive List	47
21.	Sample Strategy for Progressively Distributed Step Lists	49
22.	Comparison of Progressively Distributed and Normal Step Lists	49
23.	Comparison of Consecutive Lists With and Without Multiple Repairs per Subsystem Capability	50

LIST OF ILLUSTRATIONS (Continued)

<u>Figure</u>	<u>Title</u>	<u>Page</u>
24.	Comparison of Progressively Distributed Step and Random Spare Lists	52
25.	Comparison of Random List (per Subsystem) and Progressively Distributed Step Lists	53
26.	Comparison of Minimum and Maximum Failure Masking Lists (Order-Four Redundancy)	55
27.	Sample Strategies for Order-Two-and-One-Half Redundancy Systems.	56
28.	Comparison of Three, Order-Two-and-One-Half Failure Responsive Systems with a Third-Order Redundancy Multiple-Line System	56
29.	Comparison of Minimum and Maximum Failure Masking Lists (Order-Three-and-One-Half Redundancy)	57
30.	Order-Two-and-One-Half Progressively Distributed Step List	60
31.	Order-Three Progressively Distributed Step List	61
32.	Order-Three-and-One-Half Progressively Distributed Step List	61
33.	Order-Four Progressively Distributed Step List	62

I. INTRODUCTION

A. The Need for High System Reliability

Electronic digital data processing systems have become an integral part of the modern world. These systems are commonly used to perform tasks which were thought unachievable only a decade ago. The great computational capabilities and operating speeds of today's data processors have usually been obtained at the cost of extremely high equipment complexity. This complexity naturally results in low system reliability. This, in turn, limits the usefulness of the equipment to the extent that a paradoxical situation threatens to emerge in which system capability is extremely high but it is almost never available for use.

In addition to the problems caused by loss of operating time, high system complexity and the necessity of frequent complicated repairs aggravate the problems of supplying spare parts and properly trained maintenance personnel. These problems become increasingly troublesome as large systems are put into use at remote locations. The natural environments for most military field and shipboard equipment are sufficiently remote to make the liaison problems dominate over almost all other considerations. The limit in this area is reached by spaceborne equipment where liaison becomes virtually impossible.

The necessity for high system reliability may also be dictated by the vital nature of the system functions as well as by an interest in maximizing system usefulness or minimizing liaison problems. Quite often control systems, for example, are relatively simply in comparison to large scale data processing systems, but their continuous operation may be of vital necessity for the safety and security of an individual or a nation. The list of applications of this class includes space vehicle "on-board" controls systems,

atomic reactor controls, missile guidance and destruct systems, and secure communications systems.

B. Methods of Increasing System Reliability

1. Conservative Design

One of the first methods that design engineers successfully used to increase system reliability was that of derating electronic components. Using this procedure, circuits are designed with components of much greater power and voltage rating than the specific circuit applications require. In operation, these components are subject to such low thermal and electrical stress that their expected life approaches "shelf-life". This method has proved to be a relatively cheap and effective means for increasing average system life.

2. Hyper-reliable Components

A second method, which has been equally successful, involves the use of special manufacturing procedures to produce more reliable components. This method employs refined fabrication techniques and a supplementary program for individual component testing. The testing program is used to monitor various characteristics of the components during the manufacturing procedure such that any defects can be detected before the product reaches the consumer. This approach to achieving high system reliability has been championed by the Air Force's Minuteman Missile program. Although significant reductions in component failure rates have been realized through the use of this technique, the effort appears to be reaching a point of diminishing returns where each level of improvement is becoming more and more costly to achieve.

3. Coding

An entirely different approach to problem of achieving high reliability has been found in the use of coded signals. This approach is useful in binary data transmission and storage systems where the primary interest is that of maintaining the accuracy of existing information. In using this technique, the information to be transmitted or stored is broken up into sections called "words". Each of the words is subsequently analyzed to determine one or more of its characteristics. For example, a characteristic which is commonly of interest is the number of ones appearing in the binary word. The results of the analysis are converted to binary data, and this latter data is then combined with the original word to form a complete message unit. Depending on the complexity of the code, single or multiple error detection or correction can be performed when the message unit is decoded following transmission or storage.

In general, this technique is not applicable to systems which perform any function other than data transmission or storage. This limitation exists because any arithmetic or similar function destroys the integrity of the code by altering the message units.

4. Redundant Equipment

Several methods for achieving high system reliability through the use of redundant equipment have also been used. One relatively simple technique has been used for decades in the form of stand-by facilities. Using this method, auxiliary equipment is switched into use in the event of primary equipment failure. Most implementations of this method are extremely costly relative to the failure protection which they provide. For example, one unmaintained primary system and an unmaintained duplicate standby can only absorb one failure in each system before they both become inoperative and the

and the system function is lost. Using more sophisticated techniques, however, it is not unreasonable to expect that equipment which is replicated three or four times might absorb several dozen failures before the system function is lost.

The following section describes the basic types of redundancy techniques which have been developed. The more troublesome disadvantages of these techniques are included to provide a basis for the study reported in the remainder of this thesis.

C. Redundancy Techniques

The new techniques which have been developed for systematically introducing redundant equipment into data processing systems can be separated into two general classes: (1) component replication; (2) subsystem replication. It has been shown that the redundant equipment employed in a fixed system configuration is most effective when the system is divided into the smallest divisible units. Because the individual circuit components usually represent such units, this implies that component redundancy is the most efficient technique which can be employed. In attempting to implement redundant systems of this type, however, several problems immediately arise which suggest that this form of redundancy is not always compatible with other system design considerations.

Component redundancy is applied by placing several replicas of an electronic component in a series or a parallel configuration or some combination of the two. Examples of each type configuration is shown in figure 1. These configurations are often much more reliable than a single non-redundant component because more than one component must fail into its detrimental mode (i. e. , open or short) before the circuit function of the component is completely lost, and the system fails. For example, if a certain type diode always fails to a short mode, placing two or more of them in series as shown in figure 1a will protect the circuit from failure until all of the diodes in the

chain fail. A similar protection is provided against open circuits by paralleling components (figure 1b) or against either mode through the use of quads (figure 1c) or larger Hammock Networks (figure 1d).

It is apparent that such a technique for introducing redundancy cannot be applied to components where the actual values of the components are critical to the operation of the circuit. The failure of individual components in these configurations may easily change the impedance of the network by fifty per cent. Although most digital circuits are not particularly critical to impedance changes, many types of circuit applications are sensitive to changes of this magnitude.

In applying this type of redundancy, the assumption is made that the failure of one component is virtually independent of the operation of any other components. In systems using thin film or molecular-electronic circuits, it has been found that failures of components deposited on the same inactive base or included in the same semiconductor block are highly correlated. This means that in order to achieve even a rough approximation

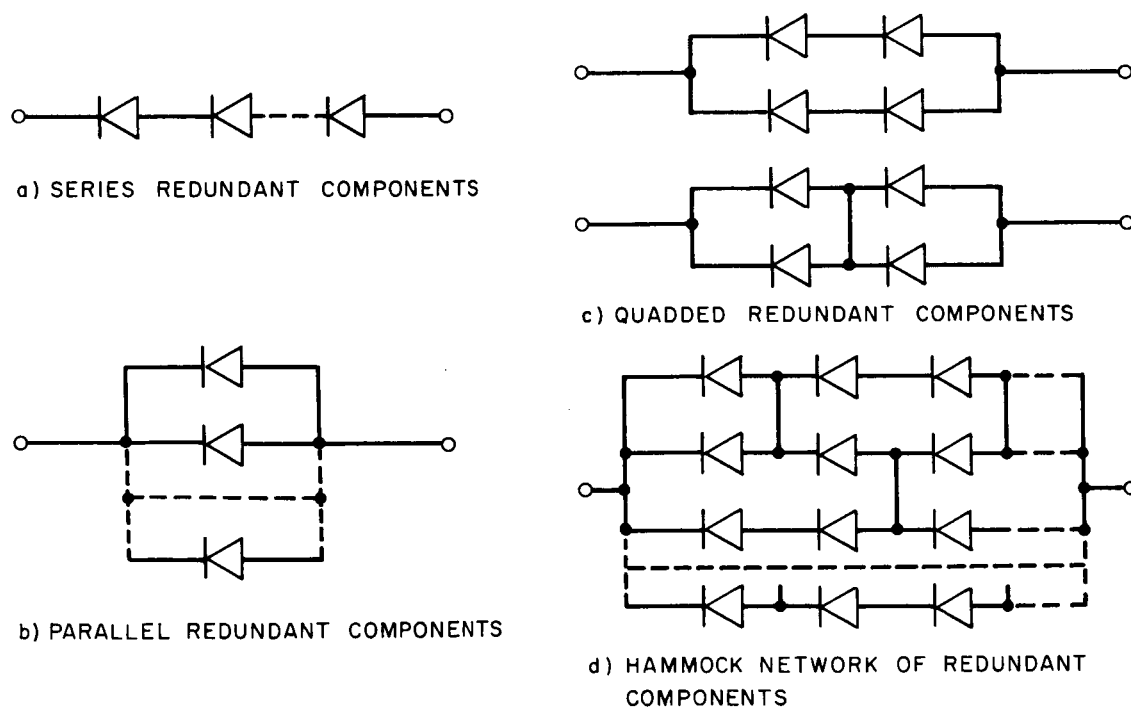


Figure 1. Redundant Component Configurations

to component independence, components in the same redundant network would have to be deposited on different bases or blocks and connected together with additional wiring. The unreliability of interconnections between these circuits would usually offset the gains sought through redundancy; therefore, a different class of techniques must be used for introducing redundancy into most microminiaturized circuits.

The second class of techniques, subsystem replication, can be subdivided into two significantly different subclasses. In the first of these, the "sense and switch" subclass, two or more nominally identical replicas of a subsystem are monitored and controlled by a monitor and control network. Based on some predetermined operational criteria the network locks the output of the stage¹ to the output of one of the subsystem replicas until a failure in that subsystem is sensed by the monitoring circuitry. At this time the control portion of the network attempts to switch the stage output to a working replica if one is available.

Although this technique is particularly useful in analog systems, it is very difficult to calculate the quality of a digital signal without comparing it to another nominally identical signal. Because of this, the sensing circuits must be very elaborate to capitalize on the advantage of one out of (n) replica operation. This is troublesome because this type operation is the major advantage derived from techniques of this subclass.

The second subclass of techniques for this type of implementation of redundant systems might be called the "voted" techniques. Of the several techniques in this subclass, the "multiple-line" method of implementation appears to be the best. Figure 2b shows basic topological characteristics of a segment of a multiple-line system. A non-redundant version of this equipment would consist of three single input, single output subsystems connected in series as shown in figure 2a. To form the redundant version

¹

A "stage" consists of all of the subsystem replicas and any associated circuitry required to provide a redundant replacement for a subsystem in a non-redundant system.

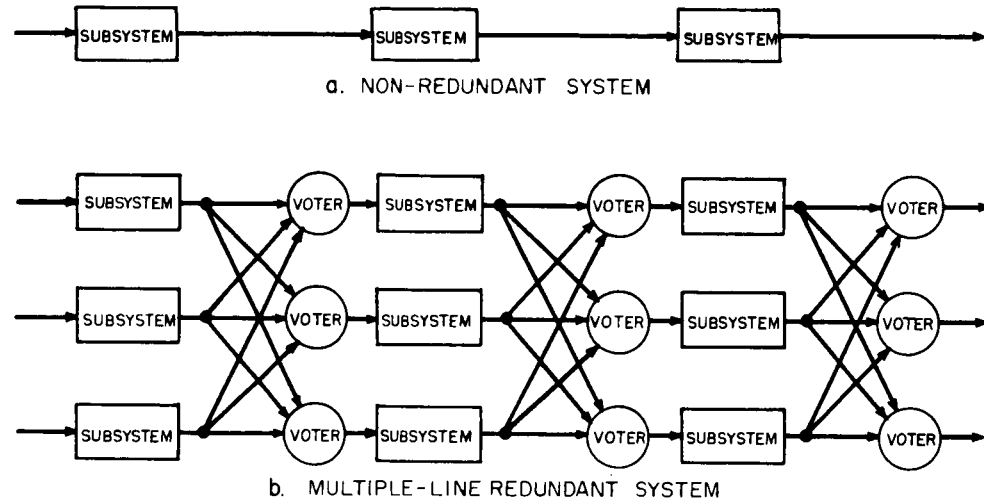


Figure 2. A segment of an Example System

of the equipment, each subsystem has been replicated twice and voting circuits (or voters) have been inserted between the sets of subsystems. The use of three subsystems to replace one from the non-redundant version results in an "order-three" system. Similarly, the use of five to replace one would result in an "order five" system. The voters are usually majority logic gates. The voters may, however, be designed to vote on some alternate threshold level. This would be done if information were available to indicate that the generation of erroneous ones is much more likely than the generation of erroneous ones is much more likely than the generation of erroneous zeros or vice versa. The replication of the voters is necessary to prevent system failure because of single failures in the voters themselves.

Several investigation teams ^{(1), (2), (3), (4)*} have studied this particular type of redundancy and found it to be applicable to a broad range of digital systems. Under the names of "Multiple-line, Majority-Voted Redundancy" and "Triple-Modular Redundancy"

* Parenthetical references placed superior to the line of text refer to the bibliography.

and possibly others, it is currently being considered by various groups for inclusion in the design of the digital portions of spaceborne equipment associated with several projects including Ranger and Saturn.

The primary disadvantage of systems of this type is that they are vulnerable to certain improbable but destructive failure patterns which may disable the system while most of the redundant equipment is still operational. One of these patterns will occur anytime two of the first few component failures happen to occur in different replicas of the same stage of an order-three system.

II. THE PURPOSE OF THIS THESIS

The techniques described above provide a variety of means for employing redundant equipment to increase the reliability of electronic digital systems. Although these techniques are effective in accomplishing the desired increases, they do not make as efficient use of the redundant equipment as would seem possible.

In this thesis, the author proposes to present the concept of a new technique which the author has developed for more efficiently using redundant equipment to increase the reliability of one class of digital systems. In addition to developing this concept, the author proposes to show that this technique is, in fact, more efficient than the comparable existing technique. The comparison of the new and the old techniques will be made through the use of results obtained from a computer simulation program which the author has developed for this specific purpose.

III. PREVIOUS WORK IN THIS AREA BY OTHER INVESTIGATORS

The use of redundant equipment has interested a relatively large number of investigators in both academic and industrial environments. The publications which have been produced by these investigators are too numerous to list here; however, a bibliography which lists over one hundred of these publications was published by P. A. Jensen ⁽⁵⁾ in 1962. The majority of this work has been concentrated on the analysis and development of fixed redundancy techniques.

Only a very few investigators seem to have seriously considered systems which are in any way similar to those of interest in this investigation. The most notable work on this latter subject appears to have been done by E. J. Kletsky ⁽⁶⁾ and S. Seshu, ⁽⁷⁾ at Syracuse University Research Institute and L. Lofgren ^{(8), (9)} at the University of Illinois Electrical Engineering Research Laboratories. Kletsky and Seshu worked as a team under a Navy contract while Lofgren simultaneously conducted an independent study for the Air Force. Both Lofgren and Kletsky were interested in developing mathematical models which would describe the expected life of systems that draw up spares from a common "pool" to perform any necessary subsystem repairs. Although Lofgren's work is generally more abstract than Kletsky's, neither of them was particularly concerned about the problems of implementing such systems. In one paper, however, Lofgren did propose a fluid flow technique for performing the subsystem replacement function. This technique is itself fraught with many problems, but it certainly represents an ingenious contribution to the art. At least one other investigator, R. R. Landers ⁽¹⁰⁾ has attempted to extend the fluid flow technique to a more nearly realizable state.

Seshu suggested two possible techniques for implementing systems of the general type that Kletsky was studying. In considering implementation, he immediately recognized the problem associated with detecting errors in systems employing a non-redundant on-line structure supplemented by a pool of spares. He proposed two feasible

implementation techniques. In one technique, he suggested that a central controller be constructed to monitor the remainder of the system and to perform any necessary subsystem replacements. As an alternative, he proposed to have a ring of subsystems with each subsystem monitoring and controlling one of its neighbors.

The system organizations described in this paper have the same general objective, i. e. , long system life, as the self-repairing systems which were considered by Kletsky, Seshu, and Lofgren. The organizational structure of the systems described here, however, are much more closely related to presently practicable digital systems than are those of the limiting cases considered by the above authors. Because of this difference between the organizational structures, this new work does not appear to be an extension of any of the other author's work.

IV. FAILURE RESPONSIVE SYSTEM ORGANIZATIONS

A. The General Concept

A "failure responsive system" is a redundant system which has the capability to partially reorganize itself to combat the detrimental effects of internal subsystem failures. Before any subsystems have failed, failure responsive systems closely resemble the multiple-line redundant systems which have been previously described. Within these systems each subsystem is also replicated several times, and each replica in each stage is supplied with a set of the inputs associated with the stage. The outputs are fed into a switching network and used to determine the best estimate of the correct stage output in a manner similar to the voting circuits of the multiple-line systems. These systems resemble the multiple-line systems until one of the stages experiences multiple subsystem failures. When this condition occurs, the switching network for that stage signals for a partial system reorganization. This reorganization consists of the elimination of the failed subsystems, and the functional movement other subsystems through the switching of their input and output connections. The result is the restoration of the system to an operational state. This process is continued as long as enough subsystems remain operational so that the reorganization action can effect the necessary restoration. It should be noted that the reorganization should not change the functional operation of the system. It only changes the distribution of the redundant subsystem replicas. As this statement implies, the subsystems which take part in the reorganizations are functionally identical so that any one can be substituted for any other one.

As an example of a typical series of operations, the reorganization actions of one system as it would respond to one particular failure pattern, will be considered. The system which will be considered is presented by the pattern of blocks shown in figure 3. This pattern of blocks represents a seven stage, order three, failure responsive system. The non-redundant version of this seven stage system would be similar to the three stage

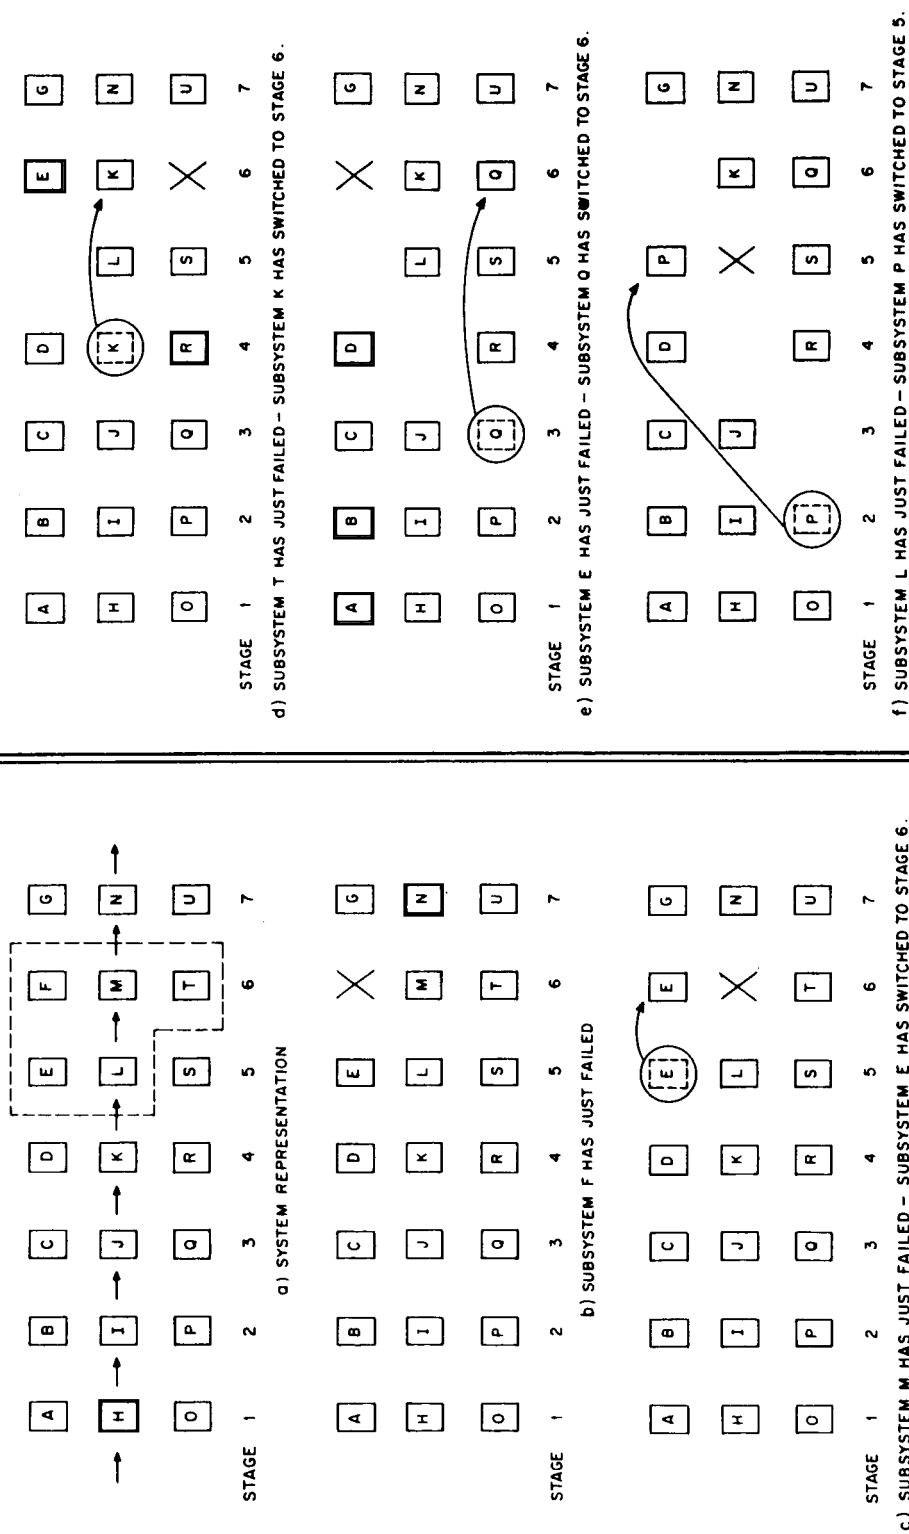


Figure 3. Example Failure Pattern

system illustrated in figure 2a. In this case and in the figures which follow, the blocks in the diagram represent individual subsystems. The physical position of the blocks represent the relative functional positions of subsystems within an electronic system. It should also be noted that the peripheral switching circuits required to implement the various systems have not been shown.

Referring to the letter code shown in figure 3a, the following series of subsystem failures are assumed to have occurred: F, M, T, E, L. Note that this series includes all of the subsystems enclosed by the dashed lines in figure 3a. Using a preprogrammed response strategy, the system would react to this pattern of failures in the following manner:

1. When F failed, its output would be permanently turned off. No other action would be taken. (See figure 3b.)
2. When M failed, the ambiguity caused by the failure of one of two nominally identical subsystems, M and T, will cause one of the working subsystems from another stage to be switched to stage 6. In this case, subsystem E will be moved up one stage. With E and T now properly performing the function of stage 6, the ambiguity existing between M and T is resolved and M is turned off. (See figure 3c.)
3. When subsystem T fails, an identical procedure will be used to call K to stage 6. Again the ambiguity existing between working subsystem E and failed subsystem T will be resolved, and T will also be turned off. (See figure 3d.)
4. When E fails, processor Q will be moved to stage 6 and again the system will be restored to operation. (See figure 3e.)
5. The subsequent failure of L, will result in subsystem P being moved to stage 5. This will restore the stage and the system to operation. (See figure 3f.)

This example was specifically chosen to illustrate the conceivable range of variation in response strategies as well as the potential power of the failure responsive technique.

The first of these two is evident. For example, although a definite response strategy was employed in the above example, it is not necessarily obvious to the reader what the strategy was, even after observing the effect of several failures in the same general location. As for the second item, it is obvious that the system would have failed after the third failure and quite probably after the second failure if the multiple-line majority voted technique were still being employed. With the failure responsive reorganization capability, however, the system has withstood five consecutive failures in a tightly grouped pattern without suffering system failure.

B. The Specific Organizational Objectives

One of the primary objectives of this study has been to develop a set of design rules for failure responsive systems. These rules are intended to serve as guidelines for facilitating system designs which will make very effective use of the redundant equipment, subject to switching network unreliability and various instantaneous failure masking requirements². To establish a meaningful set of rules, a wide variety of feasible response strategy characteristics had to be considered to determine which characteristics were necessary, which were only desirable and which were undesirable. These characteristics include the following:

1. The number of replacements which should be available to any one stage. (The assumption is made that the addition of replacements results in an addition to the peripheral switching circuitry.)
2. The pattern for specifying which subsystems should be used as the replacements for any particular stage and the order in which they should be called.
3. The use of fractional order of redundancy, i. e. not every stage being the same order in the initial state.

² "Instantaneous failure masking" means that a subsystem failure in any stage is completely masked by that stage so that no errors propagate through the system during the time the system is reorganizing itself to eliminate the failed subsystem.

4. The use of minimum order of redundancy to be maintained at a stage from which a failed stage would like to take a replacement.
5. The capability of a vulnerable stage to override the minimum of number (4) in the event no replacement is otherwise available.
6. The capability of a single subsystem to make more than one change of location.

These and other response strategy characteristics have been considered during this study. The relative importance and desirability of all of them are reflected in the conclusions presented in section IX.

V. ANALYSIS METHODS

To evaluate failure responsive systems and compare the effectiveness of various response strategies, a method had to be found for determining the reliability of these systems. In the case of multiple-line redundancy in which the functional locations are static, various analytical techniques have been used to express reliability. The problems presented in the following paragraphs indicate that the techniques used for analyzing fixed redundant systems are not generally amenable to failure responsive systems.

Before proceeding with the description of the problems involved in applying analytical techniques to failure responsive systems, it should be noted that all of the systems considered will be limited to those of simple unilateral signal flow with single inputs and single outputs at each subsystem. It is also assumed that all stages are identical; therefore, all stage reliabilities are equal. Although such systems are obviously idealistically simple, any more realistic modifications in the models would only serve to complicate the existing problem or increase the overall number of problems.

A. The "Brute Force" Method

As stated above, the assumption has been made that all stages are identical. This statement implies that the system reliability, R_s , can be expressed as

$$R_s = (R_{ST})^N \quad (1)$$

where R_{ST} = the stage reliability,

N = the number of stages in the system.

Because N is always known, the only significant problem is the determination of R_{ST} . For a system employing fixed redundancy, this problem is easily solved by enumerating the number of failure patterns which can exist within the stage and still permit stage

operation. The computation is completed by summing the probabilities that each of these patterns will exist. For example, the reliability of a stage in an order-three, majority-voted multiple line system is given by:

$$R = (e^{-\lambda t})^3 + 3 (e^{-\lambda t})^2 (1 - e^{-\lambda t}) \quad (2)^3$$

This problem is not so easily solved in the case of failure responsive systems. The mobility of the subsystems in these systems suggests that the enumeration of operating states must be performed on a complete system basis rather than be restricted to an individual stage. This approach is complicated by the fact that many response strategies are sensitive to the order in which failures occur as well as the particular locations of the failures. The number of possible operating states and the permutations of failure orders combine to make the overall reliability computation process too lengthy for practical use.

B. The Markov Chain Method

The changes in system operating states caused by subsystem failures may be regarded as transitions between states in a Markov chain. The formulation of this reliability analysis problem as a Markov chain automatically provides a group of solution methods which are not otherwise available.

Before proceeding with the analysis, however, it would seem wise to consider the size of the Markov transition matrix which would be required for the systems of interest. A typical system might have as many as a hundred or more stages in it, but to be conservative a ten-stage system will be used as an example. The number of possible operational states of a ten stage order three system is 2^{30} or 1,073,741,824. This assumes that each of the 30 subsystems is either working correctly or catastrophically failed.

³ The assumption is made that cancelling errors do not occur and the voting circuitry is perfectly reliable.

The number of entries in the transitional matrix for this system would be $(2^{30})^2$. It is obvious at this point, that even if special techniques could be found to eliminate 95% of these entries from consideration, the matrix would still be too big to handle conveniently, even using a large, high speed computer to perform the computations.

C. The Minimal Cuts Techniques

A technique for determining the lower bound on the reliability of redundant systems has been developed by Esary and Proschan⁽¹¹⁾. This technique depends on the existence of "coherent" systems and definable sets of "minimal cuts". These terms have been precisely defined by Esary and Proschan in the following manner: A system is "coherent" when it fulfills the following conditions:

- (1) A system which has failed because of a pattern of component failures existing within the system would not begin working again upon the occurrence of any additional failures.
- (2) A system which is working in the presence of a set of component failures should not stop working if any of the failed components is repaired or replaced.
- (3) A system should work when all of its components are working.
- (4) A system should fail when all of its components are failed.

A "cut" is a set of components whose simultaneous failures are sufficient to cause system failure regardless of the operational state of the other system components. (A system will usually contain a relatively large number of cuts with many components appearing in more than one cut.) A "minimal cut" is defined as any cut in which there exists no subset of components whose combined failures would cause system failure.

Failure responsive systems meet all of the conditions required of coherent systems. They do not, however, always meet the condition of definable minimal cuts. The sensitivity of many of the response strategies to the order in which failures occur destroys

the concept of a minimal cut. Figures 4a and b shows the system which illustrates this point. In the example the response strategy allows only the subsystems on the top row to change location. Any of these may move forward⁴ one or two stages if required by the existing failure pattern. If failures occur in the order indicated by the small circled numbers in figure 4a, the system will remain operational with the moveable subsystem from stage C having shifted to stage D. If however the failures occur in the order shown in figure 4b, the system fails because an unresolvable ambiguity exists in stage C. It is apparent from this example, that cuts can not always be identified by the pattern of failures existing at any particular time. This difficulty, combined with the complex problem of enumerating all the minimal cuts which can be identified, virtually prohibits the use of this analytical technique for estimating the reliability of failure responsive systems.

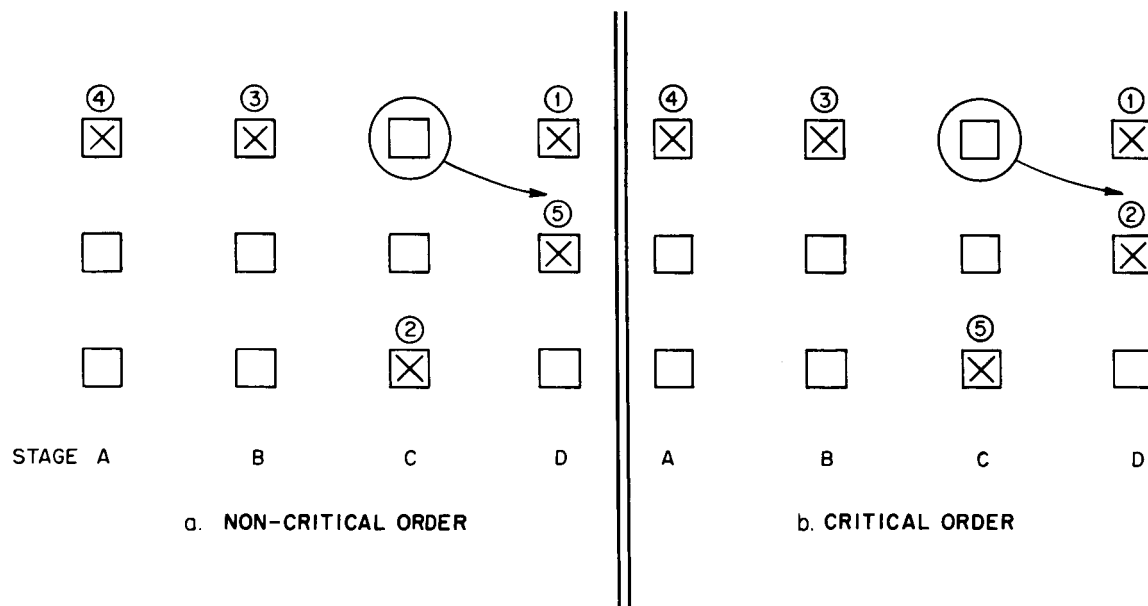


Figure 4. Critical and Non-Critical Order of Failures

⁴ Stages A and D are assumed to be adjacent so that the moveable subsystem in stage D, for example, can be moved to stages A or B.

D. The Computer Simulation Method

The concept of physically modeling a large system and testing the response of the model to gain knowledge of the true situation is a form of simulation that has been used for centuries. Although in computer simulation no physical model is built, a functional representation of a system to be tested is formed by a sequence of program statements. These statements are used to specify all of the individual deterministic actions of the system. Inputs and outputs to this model are presented to the computer in the form of data rather than physical quantities. The response of the true system to various perturbations in the input data is estimated by observing the response of the computer representation just as if a physical model had been built.

Mathematicians almost always object to the use of either physical or computer simulation because no rigorous proofs of the results can be given, and the system response cannot be described by a group of neat, closed-form expressions. Because these are valid objections, simulation analysis is usually used only for treating very large complex systems where the number of variables in the problem prohibits the use of more standard mathematical modeling techniques, or where the cost of exercising the real system is too high. In the case of failure responsive systems, the variety of characteristics inherent in the response strategies are difficult to model accurately in a mathematical expression. However, such systems can be easily handled by a computer simulation program.

The inputs to this particular program are in the form of response strategy constants, subsystem failure rates and random numbers. The random numbers are correlated with individual subsystems to represent random failures. After the simulation of several hundred input failure patterns, the program output is used to estimate system reliability.

VI. THE COMPUTER SIMULATION PROGRAM

A. The Operational Principles of the Program

The topography of a system is modeled in the computer simulation program by an array of stored data. These data can be roughly divided into two sets. The first set contains information which specifies the operating state or the characteristics of individual subsystems. The second set contains information which determines the characteristics of the overall system operation. Different system response strategies and other operational requirements are simulated by establishing, within the computer memory, the appropriate initial values of each of the stored data words. In some cases these values are read directly into the computer from an external source, while in other cases, the data is generated by the computer operating under the command of special input control constraints.

B. Individual Subsystem Information

1. Failure Location Intervals

The operation of the program is based on the assumption that subsystem failures can be simulated by the computer in such a manner that they represent the way in which actual failures would occur in operating systems. The main problem is to determine which subsystem should be designated as failed when a subsystem failure is assumed to have occurred at a particular time. In order to accurately represent the occurrence of a failure in an operating system, the conditional probability of a subsystem's just having failed, given that exactly one subsystem failure has just occurred, must be equal to the

same conditional probability that would apply to the subsystems in a comparable operating system. It is shown in the Appendix that this conditional probability is given by the simple expression:

$$P(i/1) = \frac{\lambda_i}{\sum_{i=1}^L \lambda_i} \quad (3)$$

where i refers to the i^{th} subsystem; λ_i is the failure rate of the i^{th} subsystem and L is the total number of subsystems which were operational before the occurrence of the present failure.

Randomly located subsystem failures are generated by the simulation program, subject to the above conditional probability, in the following manner. The conditional probability of failure associated with each subsystem is computed according to equation (3). The interval of numbers between zero (0) and one (1) is then divided into L subintervals with the length of each subinterval being directly proportional to conditional probability of failure of one subsystem. The assignment of one subinterval to each subsystem results in the unique association of every number in the zero (0) to one (1) interval with exactly one subsystem. To locate a simulated failure, the computer draws a random number from a population which is uniformly distributed over this same zero (0) to one (1) range. The random number thus selected must fall into one of the subintervals associated with one of the subsystems. The computer locates this subsystem and designates it as failed.

In performing this operation, the computer first reads in the failure rates of the subsystems. It then uses the failure rates to determine the conditional probabilities of failure to be associated with the subsystems, and corresponding intervals of numbers. The upper and lower bounds on the intervals then become a part of the stored data.

2. Space Lists

One of the major differences between the response strategies is the sequence in which subsystems are called to aid the failed or, in some cases, vulnerable stages. One of the outstanding features of this computer program is the simple manner in which an almost unlimited variety of sequences can be set up.

As part of the initialization procedure, an identification number is assigned to each subsystem. The sequence of subsystems to be called to aid any particular stage is established by simply reading into the computer memory a list of identification numbers. The order of the numbers combines with their actual value to precisely specify the desired sequence.⁵ The list of identification numbers is referred to as a "spare list". (This technique also permits the testing of random response strategies by the insertion of random number spare lists.)

3. Other Stored Data

In addition to the information concerning random number interval bounds and subsystem spare lists, a variety of other information is stored in the computer memory. This information is used to specify the general characteristics of the response strategy being tested and to control many of the peripheral program operations. Figures 5a and 5b illustrate a typical example of the general strategy characteristics which are specified in this manner. In both cases, stage three has experienced one failure and stage four has experienced two failures. At this point, stage four requires aid. In both cases, the first

5 It should be noted that the program is equipped with a pattern duplicating option that permits a sample spare list to be read in for one stage and the pattern reproduced for all other stages with all "spare" subsystems coming from the same relative location.

choice of a replacement is subsystem A; the second choice is B. The stored information specifying the operation of the system in figure 5a permits A to shift to stage four, leaving stage three in a non-redundant state. In contrast, the operation of the system in figure 5b restrains the movement of A because of the previous failure in the same column and forces B to aid stage four.

An example of the peripheral program operations controlled by the remaining variables is the output format. The information which is printed out by the simulation program can be manipulated so that details of the individual simulated failure patterns are available for inspection. Conversely, the output may be restricted to a brief summary of the combined statistical results of many runs.




Figure 5. Two Response Strategies

C. The Detailed Operation of the Program


1. Data Storage

The portion of the data which concerns individual subsystem operation is organized into the format of a three dimensional matrix. This matrix closely resembles the actual form of the system being simulated because two of the dimensions correspond to the number of stages and the order of redundancy of the base system. The third dimension contains data words about the subsystems represented by the first two dimensions. Figure 6a shows one such matrix which represents the typical system shown in figure 6b. As shown in figure 6a, the first two words at each location specify the random number interval bounds associated with that subsystem location. The third word specifies the identification number of that location. The fourth word is non-zero only if the simulated subsystem initially found at that location has moved or failed. If this word is non-zero, it equals the number of moves or failures which have occurred in that column at the time the particular subsystem moved or failed. The remaining data words in each matrix location are members of the spare list, where the fifth word represents the first entry on the list, the sixth word the second entry, and so forth.


The data which is stored outside this matrix applies to the overall system or program operation. This data is simply stored as individual variable values and does not form any sort of integrated data block.

2. The Simulation Procedure

After all the initial data concerning the system operation has been inserted into the computer memory, the actual simulation phase of the program begins. Although this part



a. MATRIX REPRESENTATION



b. A TYPICAL SYSTEM

Figure 6. A Typical System and Its Matrix Representation

of the program is complicated in terms of computer instructions, it is relatively simple in principle. A series of random numbers is chosen from a population uniformly distributed between zero (0) and one (1). As each number is chosen, it is associated with one of the simulated subsystems by locating the subsystem whose random number interval contains the chosen number. The failure of the subsystem is noted by adding one (1) to the previous number of failures observed in the stage to which this subsystem belongs and storing the new number in the fourth position in the matrix at that subsystem location. In addition, the random number interval bounds are set to zero (0), thus prohibiting multiple failures of any one subsystem.

After the subsystem failure has been recognized, the computer checks to see if the stage which experienced the failure subsequently requires the aid of a replacement subsystem. If the stage still meets all of the requirements imposed by all of the related criteria, no further action is taken, and the next in the series of random numbers is selected. If the stage requires aid, the program begins searching through the subsystems whose identification numbers appear on the spare list of the previously failed or moved block in the vulnerable stage.⁶

The search is conducted by interrogating the possible spares in the order in which their identification numbers appear on this spare list and determining their availability. This continues until the "repair" is made or it is determined that the repair cannot be made. If the repair can be made, the data describing the subsystem to be moved is shifted from its initial location to the location of the previous failure in the vulnerable stage. Depending on the strategy being tested, the subsystem in its new location may lose all of its remaining repair capability; it may retain its old capability, or it may assume the

6 The only case in which aid may be required by a stage which has not previously experienced a failure or the loss of a subsystem to another stage is in systems having unequal stage redundancy. The program then considers the low order stages as having lost some subsystems.

capability of the subsystem which it replaced. If the repair cannot be made, a check is made to see if the number of operating subsystems remaining in the vulnerable stage is two or greater. If the answer is yes, the simulation continues with the drawing of another random number. If the answer is no, it is assumed that the most recent failure has resulted in the occurrence of an unresolvable ambiguity in the vulnerable stage; therefore, the system has failed.

The procedure is continued until the system reaches the failed state. At this point, the total number of subsystem failures in the system is recorded, the matrix is reset to the original state, and the entire procedure begins again. The repetition of this procedure several hundred times produces statistical information which can be used to construct estimates of the reliability versus time curves of systems using the response strategy being tested. The entire simulation procedure is summarized by the flow chart in figure 7.

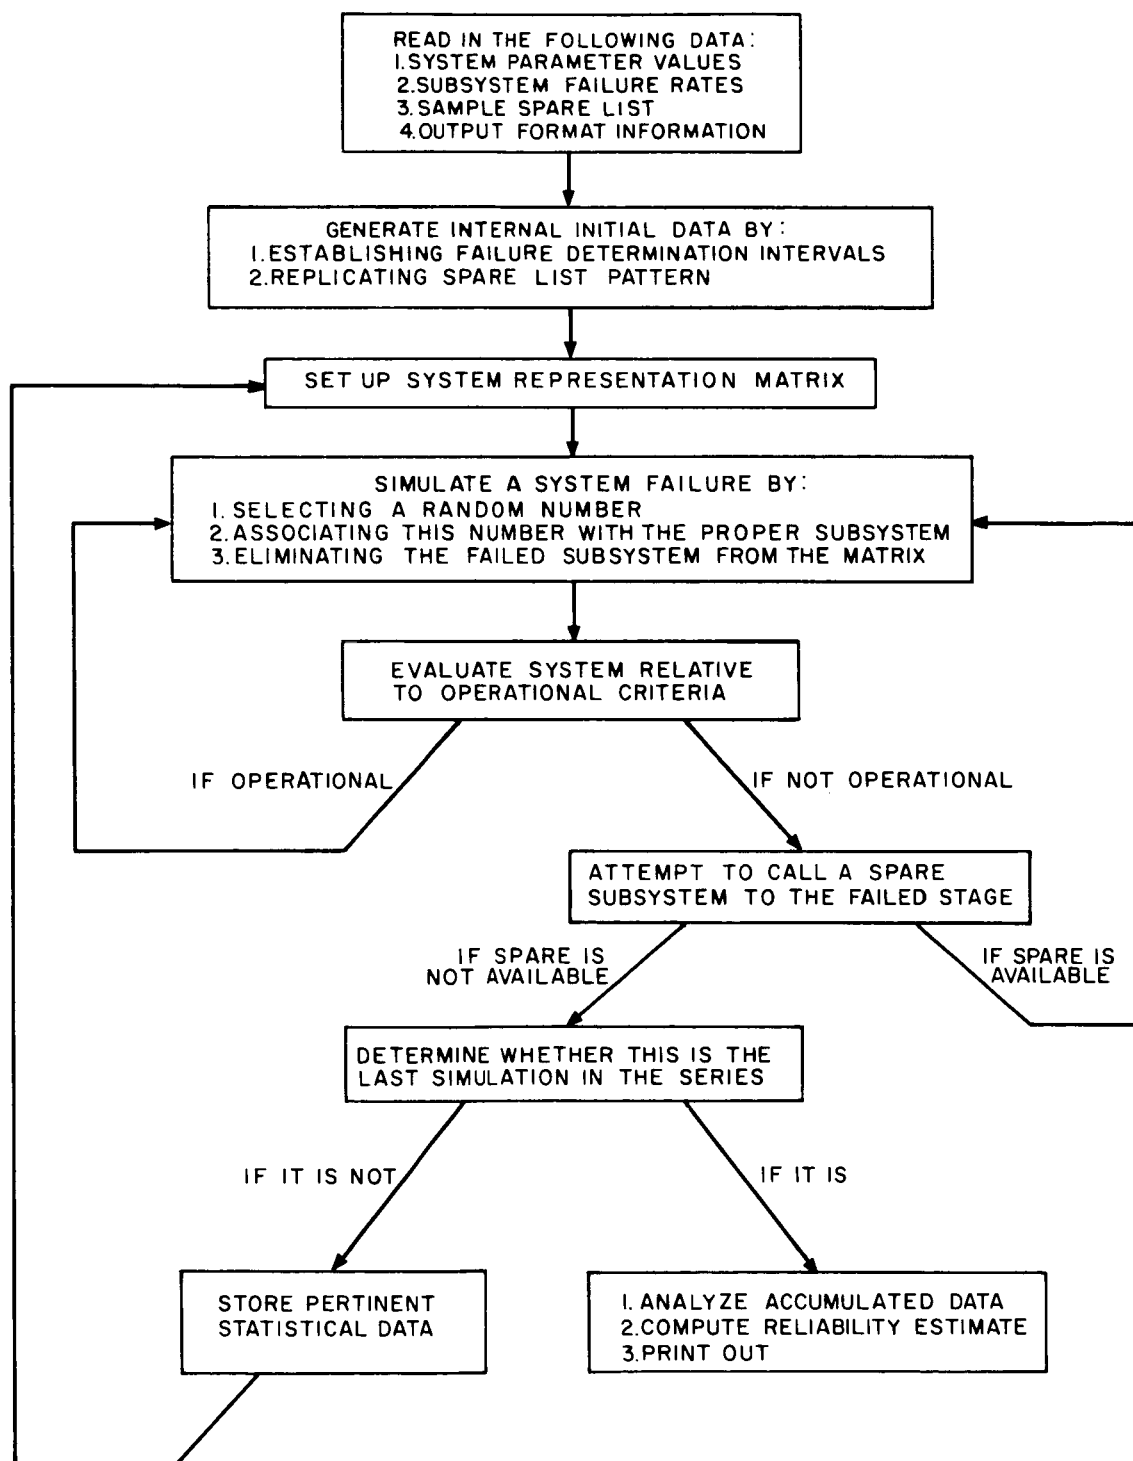


Figure 7. Summary Flow Chart of Computer Program

VII. SYSTEM EVALUATION

A. Methods For Estimating System Reliability Versus Time Curve

1. The Conditional Probability Method

The information obtained from the simulation procedure can be used to construct a histogram which describes the relative observed frequency of system failures for any given number of subsystem failures. Figure 8 shows an example of such a histogram. The height of the lines $f(x)$ in this histogram are determined by counting the number of systems which were observed to fail with exactly x subsystem failures in the system and dividing this number by the total number of system failures which were simulated. Thus, the magnitude of these lines represent a statistical estimate of the probability that a particular system will fail at the occurrence of exactly the x th subsystem failure.




Figure 8. Histogram of Observed System Failures

Figure 9 shows the cumulative curve which is formed by adding segments of the above histogram according to the relationship

$$F(x) = \sum_{i=0}^x f(x) \quad (4)$$

The magnitude of $F(x)$ is an estimate of the conditional probability that a system has failed, given that exactly x failures exist within the system. It is this probability that is needed to calculate the system reliability.

If the assumption is made that the failure rates of all the subsystems are equal, the probability of exactly x failures occurring in a system containing N subsystems can be calculated from the expression

$$P(x, t) = \binom{N}{x} (1 - e^{-\lambda t})^x (e^{-\lambda t})^{N-x} \quad (5)$$

where,

$\binom{N}{x}$ is the symbol for x combinations of N items.

This probability can be combined with the estimated conditional probability of system failure to produce an estimate of the overall system reliability. This can be done using the relationship

$$R(t) = \sum_{x=0}^N F(x) P(x, t). \quad (6)$$

To apply this technique to non-homogeneous systems having more than one subsystem failure rate, two alternative possibilities have been considered. By recording the distribution of failures among the different types of subsystems, the individual lines




Figure 9. Cumulative History of Observed System Failures

shown on the histogram could be subdivided so that their magnitudes represented the conditional probability that system had failed, given that the system has absorbed x failures of one type subsystem, y failures of another types subsystem, z failures of another and so forth. To obtain meaningful estimates of each of the conditional probabilities which can be defined in this manner, an unreasonably large total number of system failures would have to be simulated.

A much simpler method, which is equally accurate for a limited number of samples⁷, has been used in this program. In this second method a weighted average⁸ of the various subsystem failure rates is computed, and this number is substituted for the single failure

⁷ i. e. , 500 to 1000

⁸ $\sum_{i=1}^N \frac{m_i \lambda_i}{N}$ where m_i is the number of subsystem subject to the failure rate λ_i .

rate used in the equation given above for computing the reliability of homogeneous systems. It has been found experimentally that the random error introduced by the generation of random failures usually masks out completely any error introduced by the use of the weighted average.

2. The Random Time Generation Method

In addition to the simulation of random failure patterns, the computer program can be used to locate randomly in time the occurrence of each failure in a pattern. It has been previously stated that each subsystem is subject to a constant failure rate. This implies that the probability of continuous operation of all (N) subsystems in any system from the time $t=0$ is given by the expression

$$R(t) = e^{-\sum_{i=1}^N \lambda_i t} \quad (7)$$

Conversely, the probability that the first subsystem failure will occur in the interval of time zero to t is given by the expression

$$P(1^{st}) = 1 - R(t) = 1 - e^{-\sum_{i=1}^N \lambda_i t} \quad (8)$$

Using a relationship described by A. M. Mood,⁽¹²⁾ a set of random numbers drawn from a population uniformly distributed between zero and one can be transformed to a similar set of random numbers belonging to any other distribution. For the case of the exponential distribution of interest, this is accomplished by letting

$$f(y) = 1 \quad \text{for } 0 \leq y \leq 1 \quad (9)$$

$$f(y) = 0 \quad \text{elsewhere} \quad (10)$$

$$\text{and } y = G(t) = 1 - e^{-\lambda_s t} \quad \text{where } \lambda_s = \sum_{i=1}^N \lambda_i \quad (11)$$

Figure 10 shows this last relationship graphically.




Figure 10. Uniform to $G(y)$ Distribution Transformation

Using this relationship, a random number taken from a set of uniformly distributed numbers is used to generate the time to the first subsystem failure with the correct probability of picking a time from any increment along the time axis. By simply subtracting the failure rate of the first failed subsystem from the total failure rate λ_s and setting the time scale reference at the point of the first failure, a time between the first and second failure can be determined in the same fashion. The sum of these two times simulates the total system operating time up to that point.

This process is repeated until the system withstands so many failures that it fails to meet the system operational criteria. The occurrence of this event stops the procedure, and the various system state change characteristics and the total operating time are recorded.

The record of total operating times can be directly used to estimate system reliability⁹ versus time. This is done by ordering the individual operating times so that the percentage of systems operating prior to any given time can be calculated. This percentage is exactly the observed system reliability and may be used as an estimate of the true system reliability. It should be noted that the observed system reliability is always constant between observed system failure times, therefore, a discontinuous curve such as that labeled "A" in figure 11 results from the unmodified use of this estimation procedure. A much smoother curve can be obtained by interpolating intermediate values in the area between the points.

- 9 Reliability, is defined here as the probability of continuous system operation over a time interval zero (0) to (t) when it is known that the system was operating at time zero (0).




Figure 11. Comparison of Reliability Estimation Curves

3. Comparison of the Two Estimation Techniques

The reliability curves produced by both of these estimation techniques tend to be more accurate in the central region of the curves where most system failures occur than they are at either of the upper or lower extremities. This situation exists because the extreme regions are dominated by the few system failures which occur either with very few subsystem failures or unusually many subsystem failures having been withstood. No rigorous method for evaluating the two estimation techniques has been devised. For the purpose of this study, the equation method of estimation has been chosen rather than the time generation method. This choice was based on fact that the equation method required no sophisticated method of interpolating between observed points to provide meaningful estimates of the shape of the reliability curve in the high reliability region. The curves shown in figure 11 may help clarify this point. In contrast to the five events which control the shape of the step curve which naturally results from a sample of 500 events using the time generation technique, all the 500 events contribute in some amount to the continuous curve produced by the equation technique.

B. Single-Valued Measures of Performance

The techniques which have been described above provide an estimate of system reliability as a function of time. Because the comparison of the reliability of various systems at every point in time is not practical or particularly meaningful, a method of using the functional reliability estimate to generate a single-valued measure of performance had to be found. The several possibilities which have been considered are described below.

1. Mean Time Before Failures

The most popular reliability measure applied to non-redundant systems is the "mean time before failures" or "MTBF". The MTBF is a quite useful reliability measure for non-redundant systems of this type because the associated reliability curves are all of the exponential form, having time constants which are inversely proportional to the MTBF (see figure 12). This measure is not as meaningful for failure responsive systems whose reliability curves vary in form. Figure 13 shows two curves which have approximately the same MTBF's, but they are obviously not equivalent systems. It would seem, therefore, that a more useful measure should be found.

2. System Reliability at a Selected Time

The reliability of systems at one point in time is an alternate measure that deserves consideration. This is by far the easiest measure to compute, but it does have some inherent disadvantages. This measure may simply show that one system is more reliable than another system at one particular point in time. If a situation such as the one illustrated in figure 13 exists, the system which is more reliable at t_1 may not be the more desirable if the mission is completed at t_2 . Similarly, two systems may appear to be nearly equivalent at the evaluation time when they differ greatly before the end of the mission time. Figure 14 shows examples of the reliability curves of two such systems. Again, it would seem that a still better measure should be found.




Figure 12. Non-Redundant System Reliability Curves




Figure 13. Different Reliability Curves with Similar Means




Figure 14. Different System Reliability Curves with Similar Short Life Reliabilities

3. Quantile Occurrence

The third available evaluator has not been extensively used in the past, but it seems to overcome some of the disadvantages of the first two possibilities. This method uses the time at which the system reliability falls below a pre-determined quantile as the measure of evaluation. This measure is defined here as the "useful life". Figure 15 illustrates the method for the 0.90 quantile. In this case, the system characterized by the 0.90 quantile occurring at t_2 is more desirable than the system with the quantile occurring at t_1 . This evaluator tends to overcome the problem inherent in the MTBF evaluator because only the region of the curve which is of interest enters into the evaluation. The problem of performing the evaluation only at a single point in time, which is associated with the second evaluator is also solved because this third evaluator is more sensitive to differences in system reliabilities in the high reliability region.




Figure 15. The "Useful Life" Measure

The problem of accepting the wrong system because of curve crossover may be virtually eliminated by confining the quantile selected as the criteria to the high reliability region. This is not a particularly significant restraint because the nature of the applications, which require the use of the sophisticated systems being considered here, will require operation strictly in the high reliability region.

VIII. SIMULATION RESULTS

The simulation study of response strategies has been conducted in two phases. In the first phase the assumption was made that all of the peripheral error detection and switching circuitry required to implement the systems was perfectly reliable. In the second phase, this assumption was dropped and a failure rate was associated with the peripheral circuitry as it is with the functional subsystems.

Although the first phase effort may appear to be completely superfluous when compared to the second phase, this is not the case in practice. The first phase results indicate which response strategies are optimal if it is given that certain numbers of subsystems appear on the individual spare lists. This optimal strategy information is independent of the failure rate of the switching circuitry. The second phase results merely show what the length of the spare list should be, given the failure rate of the subsystems, the minimum peripheral circuitry failure rate, and the additional failure rate which must be added to the minimum to account for each addition to a spare list.

In the pages which follow, the results which have been obtained during both phases of the study are described. To obtain each point estimate of the reliability of systems using any of the response strategies, the simulated systems have been subjected to five hundred sets of failure patterns of sequentially generated subsystem failures. The patterns contain the minimum number of subsystem failures required to cause system failure when the subsystem failures occur in the order generated.

The curves shown below were constructed by plotting the time of occurrence of the 0.90 quantile on the estimated reliability curves. All the curves represent systems of twenty stages, with the subsystem failure rate constant for all stages in all systems. The original order of redundancy of the systems tested is noted in the subsection title.

A. Phase I Simulations

1. Order-Three Systems

a. Experiment I. In the first set of response strategies to be tested, the capability of a subsystem to serve as a spare (or a replacement) was restricted to one subsystem in each stage. The difference between the strategies stemmed from the pattern and the order in which the subsystems having the spare capability appeared on the spare lists of the individual stages. Three subsets of strategies were tested in the course of this experiment. Figure 16 shows a sample spare list for one stage of each subset. The spare list pattern is replicated for each stage, with the first and last stages assumed to be adjacent, thus forming a closed "loop". The members of each subset all employ the same spare list pattern. The individual members of a subset may be distinguished from each other by the number of subsystems composing their associated spare lists.

The object of this experiment was two fold. The first objective was to attempt to verify the null hypothesis that the individual strategies were pair-wise equivalent, i. e. that only the length of the spare lists was significant, and not the selection pattern. The second objective was to determine the effect of allowing systems to have a "rescan" capability. A system with rescan capability is one which first scans a spare list attempting to find and call up a replacement subsystem only from a stage which has experienced no failures. If no replacements are found, it will "rescan" the list, searching for a subsystem from any stage which has more than one operating subsystem.

Figures 17 and 18 shows the results of this experiment graphically by the curves. It is apparent from these curves that the difference between spare list patterns (i. e. , sequential, uniformly distributed, or alternating consecutive spare lists) is insignificant, but that the rescan capability does have a significant effect.






Figure 16. Sample Strategies for Consecutive Lists




Figure 17. Comparison of Alternating and Sequential Consecutive Lists

b. Experiment II. In the second set of strategies to be tested, the single spare per stage restriction was released and any subsystem in a stage was allowed to perform as a spare if the spare list lengths required. Figure 19 shows the "normal step" pattern which was the basic pattern used for all the strategies in this class. The only difference in the strategies was the length of the spare lists. The object of this experiment was to determine the effect produced by spreading the spare capability among more subsystems with less movement capability.




Figure 18. Comparison of Response Strategies with and without Rescan Capability




Figure 19. Sample Strategy for a Normal Step List

Curve 1 in figure 20 shows the results of this experiment relative to curve 2, the curve for the consecutive lists from Experiment I. It can be seen from these curves that the use of the step list results in a pronounced improvement over the consecutive list system.

c. Experiment III. The next set of response strategies to be tested can be described as modifications of the step list strategies tested in Experiment II. Figure 21 shows an example spare list pattern used by these strategies. The close resemblance to the step list pattern is immediately apparent. The primary difference between the two sets of strategies is the distribution of the stages from which the spares are drawn. The strategies tested in this experiment tend to reduce the mutual dependence of any two stages on replacement




Figure 20. Comparison of Normal Step and Consecutive Lists

subsystems from the same stages. The object of this experiment was to test the effect of this reduction in the mutual dependence.

Figure 22 shows the curves which indicate the effect achieved by progressively distributing the spares. The relatively minor gains which are made by this simple modification are significant, however, because they can be achieved without increasing the amount of peripheral circuitry, regardless of the type circuitry which is used.

d. Experiment IV. All of the strategies considered in the first three experiments have restricted spare subsystems to making only one of its possible moves. Thus, if a subsystem moved to a new location and made a repair, every spare list on which that subsystem originally appeared was effectively shortened by one entry. The set of strategies which were tested in this experiment employed spare lists which were identical to those of the consecutive and distributed step list used previously. The only difference was that subsystems were allowed to move to the aid of vulnerable stages without regard to whether they had moved previously. The object of this experiment was to determine if this "multiple-move" capability would be significant in improving system reliability.

Figure 23 shows the results of the simulation graphically. Again, one of the consecutive list curves developed in the earlier experiments is included in figure 23 to provide a reference for the degree of improvement. It can be seen from this figure that a slight improvement is obtained through the addition of the multiple-move capability, but it is not nearly as pronounced as some of the other effects have been.

This same experiment was conducted using the progressively distributed step list. In this case, the curves were precisely the same for systems having less than four spares on spare lists of the individual stages. For systems having four or more spares, the curves were so nearly the same that the difference could not be observed from plots made to the same scale as the rest of the curves presented in this paper. In retrospect, this




Figure 21. Sample Strategy for Progressively Distributed Step Lists




Figure 22. Comparison of Progressively Distributed and Normal Step Lists




Figure 23. Comparison of Consecutive Lists With and Without Multiple Repairs Per Subsystem Capability

result could have been at least partially anticipated because the subsystems in systems having less than four spares per stage have only one movement possibility; therefore, multiple moves are inherently impossible. For systems having four or more spares per stage, the chance of requiring multiple moves is apparently very low.

e. Experiment V. Each of the first four experiments was designed to test the effect of some particular characteristic of systems using well-ordered response strategies. The response strategies which were simulated in this experiment do not belong to this well-ordered class. The spare lists for this set of strategies were, in fact, completely random.

Random patterns for each stage were generated by forming the spare lists from a set of randomly selected identification numbers. With the exception of those subsystems originally located in the stage for which the spare list was being generated, all of the I. D. numbers in the system were available each time a selection was made.

The primary object of the experiment was to test the relative effectiveness of the well-ordered strategies by determining the reliability of a system using different randomly selected spare lists for each stage. Figure 24 shows the results obtained from this experiment. It can be seen from the comparison of curve (1) with curve (2) in figure 24, that the random strategy is certainly not as bad as might be suspected. This is true because the mutual dependence of any two stages on spares from any other stage is relatively low, and the spare capability is spread among all the subsystems in the system. As it was shown in Experiment III, these two factors are very effective in improving system reliability. Furthermore, it should be noted that the results shown in figure 24 correspond to a random system which was found to be the best of several such systems tested.

As a matter of interest to the investigator, another set of random strategies was simulated. This set was permitted to have a different spare list for each individual subsystem. The only restriction which was imposed was that no subsystem spare list could include the identification numbers of subsystems located in the same stage as the subsystem for which the list was being prepared.

The object of this portion of the experiment was to determine if systems using individual subsystem spare lists were potentially more reliable than those which are restricted to one list per stage. Figure 25 shows the results of the simulation. Although the results shown here do indicate that such systems offer a slight advantage in the lower region of the curve, the investigator judged the implementation problems of this type response strategy to be too formidable to merit further study at this time.




Figure 24. Comparison of Progressively Distributed Step and Random Spare Lists

2. Order-Four Systems (Experiment VI)

Higher order redundant systems may be used to reach either of two objectives. One of these objectives is the achievement of longer system life through the provision of additional failure absorption capability. The second is achievement of a high degree of instantaneous failure masking capability. There is a relationship between these two objectives which inherently results in the partial realization of both effects whenever one is sought. There is, however, a definite difference between the system structure required to maximize either effect. In the long life case, the systems are organized so




Figure 25. Comparison of Random List (Per Subsystem) and Progressively Distributed Step Lists

that no repairs are effected until a stage has experienced a subsystem failure which causes an unresolvable ambiguity to exist in that stage. This is the same switching criteria used for the order-three systems. In the high failure masking case, repairs are performed whenever a subsystem failure results in less than order-three redundancy being maintained at any stage. It should be noted that the assumption has been made in both cases that any subsystem may move only one time, i. e., may make only one repair.

Based on the preceding test results, only the progressively distributed step list response strategy was considered. The simulation tests for the order four systems were used to determine the relative potential difference between systems subject to different

failure masking restraints. Figure 26 shows the results of the test. As might be expected, the early use of spares to provide instantaneous failure masking capability precludes their later use for greatly extending the life of a system after it has experienced a relatively large number of failures.

3. Fractional Order Systems

a. Experiment VII. The serious consideration of less than order three redundancy for systems using the multiple-line configuration is virtually impossible. Certainly no consideration would be given to making any stages second order because these stages would be twice as vulnerable to failure as their non-redundant counterparts. Systems of this type are, however, quite practicable when the systems have some failure responsive capability.

Figure 27 shows two, "two-and-one-half" order system. As the figure illustrates, these systems have third-order redundancy at half their stages and second order at the other half. The use of fractional order system introduces some interesting new problems. For example, if consecutive lists are to be considered, the problem of where to put the "empty spots" in the system immediately arises. Figures 27a and 27b illustrate the two must divergent possibilities. Figure 27a schematically shows a system having the "empty spots" in the row from which spares are taken. Figure 27b shows a similar system having the "empty spots" in a different row. Figure 28 shows the curves which compare these two possibilities and the progressively distributed step list. The most significant item to be found in figure 28 is the potential improvement in useful system life over the order-three multiple-line configuration by failure responsive systems having less than order-three redundancy.




Figure 26. Comparison of Minimum and Maximum Failure Masking Lists (Order-Four Redundancy)




Figure 27. Sample Strategies for Order-Two-and-One-Half Redundancy Systems




Figure 28. Comparison of Three, Order-Two-and-One-Half Failure Responsive Systems With a Third-Order Redundancy Multiple-Line System

b. Experiment VIII. In the same manner that systems can be designed using two-and-one-half-order redundancy, they can be designed using three and one half order redundancy. The primary reasons for employing this greater order of redundancy are identical to those associated with the order-four systems, i. e. , longer life or higher instantaneous failure masking capability. As in the case of the order-four systems, the achievement of high instantaneous failure masking results in a shorter overall "useful" life. It is important to note, however, that even under the high degree of failure masking restraint, these systems have potentially much longer lives than either order three systems or fixed redundant (i.e., no spares) order-three-and-one-half systems. Figure 29 shows the curves illustrating all of these effects.




Figure 29. Comparison of Minimum and Maximum Failure Masking Lists
(Order-Three-and-One-Half Redundancy)

B. Phase II Simulations

In all of the experiments which were conducted during Phase I of this study, the assumption was made that failures could not occur in the peripheral switching circuitry required to implement the response strategies. For the experiments of Phase II, this assumption has been dropped and a much less restrictive, more realistic set of three assumptions has been substituted. These assumptions may be stated as follows:

1. All the peripheral error detection switching circuitry may be divided into sections which can be uniquely associated with a single subsystem.
2. The failure of a detection and switching circuit will have the same effect as the failure of the associated subsystem.
3. The error detection and switching circuitry associated with any particular subsystem may be subdivided into a fixed portion (FSC) and a variable portion (VSC). The fixed portion represents the minimum amount of circuitry required by the subsystem to operate in its original location. The variable portion is the amount of added circuitry required by the subsystem to move to each new location.

The relative failure rates of the subsystems, the FSC and the VSC are represented in the following discussion and figures by the designations:

$$\text{Subsystem Failure Rate} = \lambda_{SS} \quad (12)$$

$$\text{FSC Failure Rate} = \lambda_{FSC} \quad (13)$$

$$\text{VSC Failure Rate} = \lambda_{VSC} \quad (14)$$

In this study, only the relative failure rates were of interest; therefore, these rates are expressed in "units", rather than in parts per hour or any other specific units.

An example of how these relative failure rates are used to compute the total relative failure rates of individual subsystems is given below. For this example, the following assumptions are made:

1. The relative failure rates are:

$$\lambda_{SS} = 1.0 \text{ Units} \quad (15)$$

$$\lambda_{FSC} = 0.2 \text{ Units} \quad (16)$$

$$\lambda_{VSC} = 0.5 \text{ Units} \quad (17)$$

2. An order-three redundancy system with four spares per stage and a progressively distributed step list is being considered. (This assumption means that two thirds of the subsystems in the system will have the capability to move to one new location and the remaining third can move to two locations.)

The total relative failure rate of the subsystems which can move to one new location is:

$$\lambda_{SS} = 1.0 \quad (18)$$

$$\lambda_{FSC} = 0.2 \quad (19)$$

$$\lambda_{VSC} = 0.5 \quad (20)$$

$$\lambda_{TOT} = 1.7 \text{ Units} \quad (21)$$

The total relative failure rate of the subsystems which can move to two new locations is:

$$\lambda_{SS} = 1.0 \quad (22)$$

$$\lambda_{FSC} = 0.2 \quad (23)$$

$$\lambda_{VSC} = 2 \times 0.5 = 1.0 \quad (24)$$

$$TOT = 2.2 \text{ Units} \quad (25)$$

These total failure rates may be interpreted to mean that the switching circuitry associated with a particular subsystem is approximately 0.70 or 1.20 times as "complex" as the

subsystem, respectively. The results of the experiments conducted during Phase I of this program indicate that the progressively distributed step list response strategy is generally the most effective of the strategies considered. For this reason, the experiments of Phase II have been limited to systems using the progressively distributed step list response strategy.

The objective of these experiments was to show that the addition of failure responsive capability would be highly beneficial to redundant system life even if the error detection and switching circuitry were relatively unreliable. To accomplish this, the relative failure rates used in the above example were applied to the order two and one half, order three, order three and one half and order four systems. Figures 30, 31, 32, and 33 show these results.




Figure 30. Order-Two-and-One-Half Progressively Distributed Step List




Figure 31. Order-Three Progressively Distributed Step List




Figure 32. Order-Three-and-One-Half Progressively Distributed Step List




Figure 33. Order-Four Progressively Distributed Step List

IX. SUMMARY AND CONCLUSIONS

A. Summary

The primary objective of this study has been the development of a new technique for more effectively employing redundant equipment to increase the useful life of electronic digital systems. Such a technique has been devised for the class of digital systems having a high degree of homogeneity among the subsystems within each system. This thesis describes the work by the author in developing this technique and in evaluating its effect upon the reliability of this particular class of digital systems.

The sections of this thesis can be divided into three groups. The first group indicates the need for highly reliable systems and describes a few of the techniques which have been developed for achieving high reliability. This group includes a description of the failure responsive systems whose characteristics are of primary interest in this investigation.

The material presented in the second group describes several techniques which were considered in attempts to develop mathematical expressions for the reliability of failure responsive systems. The failure of the techniques to describe adequately these systems resulted in the formation of a computer simulation program. The details of this program are presented in Section VI. The final portion of this group describes the measure of effectiveness which was established as a means for comparing the different organizational strategies discussed in the thesis.

The last group contains a description of the results which have been obtained from the simulation program. The curves presented in Section VIII represent the combined results of thousands of simulated system failures. The conclusions which can be reached from observing the curves of Section VIII are listed in this final section.

B. Conclusions

The curves (figures 30, 31, 32, and 33) presented in Section VIII of this thesis show that the progressively distributed step list response strategies are the most efficient of all the well-ordered strategies which were tested. The observance of this characteristic and the recognition of the value of "rescan" capability leads to the following general conclusions:

1. The capability of individual subsystems to move to new locations should be as evenly distributed among the subsystems as possible.
2. The subsystems which are available for use as spare (or replacements) to any two stages should be chosen so that the mutual dependence by these stages on the same spares is minimized.
3. The systems should be so organized that, in normal circumstances, a subsystem will not move to the aid of a critically failed stage if its movement will leave the stage in which it is presently operating vulnerable to a single failure. A critically failed stage should have the "authority", however, to demand the movement of a spare subsystem if the movement of all of the spare subsystems available to this stage are restricted as above.

It can also be concluded that order-two-and-one-half redundant failure responsive systems may effectively replace order three redundant multiple-line systems in applications where instantaneous failure masking is not important. Conversely, applications with either high instantaneous failure masking or exceptionally long life requirements may be benefitted employing order-three-and-one-half or order-four redundant failure responsive system to replace order-three, or even order-five, multiple-line systems.

From figures 30, 31, 32, and 33 presented in Section VIII, it may be concluded that the beneficial effects obtained from failure responsive capability more than offsets the disadvantages inherent in the relatively complicated circuitry required for system implementation. These curves show that the useful lives of the example systems have been significantly increased over those of the corresponding examples of multiple-line systems. These increases have been realized despite relatively pessimistic assumptions regarding the reliability of the error detection and switching circuitry.

Finally, it may be concluded that the optimum number of spare subsystems which should be made available to any stage is a function of the failure rate of the peripheral circuitry relative to the failure rate of the subsystems. It can be seen from the curves in figures 30 through 33 that for systems having relatively simple subsystems the optimum number of available spare subsystems per stage will be around three to five.

Based on all of the above, the general conclusion may be drawn that failure responsive systems do employ redundant equipment more effectively than the fixed redundant systems previously developed. The requirement of homogeneous subsystems limits the usefulness of the failure responsive technique, however, because only a relatively small class of digital systems has this homogeneous characteristic.

BIBLIOGRAPHY

1. Kemp, John C. , "Redundant Digital Systems, " Redundancy Techniques for Compu-
tin Systems, Spartan Books, Washington, D. C. , February 1962.
2. Mann, W. C. "Systematically Introduced Redundancy in Logical Systems, " 1961
IRE International Conv. Rec. , 9, P + 2, March 1961.
3. McReynolds, J. , "Evaluation of the Majority Principle as a Technique for Improving
Digital System Reliability", Hycon Eastern Inc. , (now Hermes Electronics, a
Division of Itek), Cambridge, Mass. , July 8, 1958.
4. Ramer, Paul and Carlo Michel, "Improved System Reliability by Means of Equip-
ment Redundancy, " Electronic Systems and Products Division, Martin Marietta
Corp. , October 1963.
5. Jensen, P. A. , "Bibliography on Redundancy Techniques", Redundancy Techniques
for Computing Systems, Spartan Books, Washington, D.C. , February 1962.
6. Kletsky, E. J. , "Self-Repairing Machines, " Final Report Part One (RADC-TR-61-
01B), Syracuse University Research Institute, Syracuse, New York, April 1961.
7. Seshu, Sundaram, "Self-Repairing Machines", Final Report Part Two (RADC-TR-
91B), Syracuse University Research Institute, Syracuse, New York, April 1961.
8. Lofgren, Lars, "Qualitative Limits for Automatic Error Correction Self-Repair, "
Tech. Report 6, University of Illinois Electrical Engineering Research Laboratories,
Urbana, Illinois, June 1960.
9. Lofgren, Lars, "Kinematic and Tessellation Models of Self-Repair, " Tech. Report
8, University of Illinois Electrical Engineering Research Laboratories, Urbana,
Illinois, December 1961.

10. Landers, R. R. , "Machines That Grow", Machine Design Vol. 34, No. 16, July 6, 1962.
11. Esary, J. D. and F. Proschan, "The Reliability of Coherent Systems," Redundancy Techniques for Computing Systems, Ed. by R. H. Wilcox and W. C. Mann, Spartan Books, Washington, D. C. 1962 (pp. 47-61).
12. Mood, A. M. , Introduction to the Theory of Statistics, McGraw-Hill Book Co., Inc. , New York, 1950, page 107.
13. Sasieni, Maurice, et al, Operations Research Methods and Problems, John Wiley and Sons, Inc. , New York, 1961, page 126.

APPENDIX

The assumption has been made in this thesis that individual subsystems fail at random times but at some constant rate, lamda (λ). The fact that the rate is independent of time implies that the probability of failure (13) of any one subsystem in any interval Δt is

$$\text{Pr (failure)} = \lambda \Delta t \quad (26)$$

if the interval is sufficiently small. If a subsystem failure is known to have occurred in some interval Δt about the time t , and one is interested in the conditional probability that the failure occurred in one of a set of identical subsystems, the following relationship can be seen to exist:

$$P (\text{failure of the } i\text{th subsystem/one subsystem has failed}) = \frac{\lambda_i \Delta t}{\sum_{\text{All } i} \lambda_i \Delta t} \quad (27)$$

$$\text{but} \quad \frac{\lambda_i \Delta t}{\sum_{\text{All } i} \lambda_i \Delta t} = \frac{\lambda_i \Delta t}{\Delta t \sum_{\text{All } i} \lambda_i} = \frac{\lambda_i}{\sum_{\text{All } i} \lambda_i} \quad (28)$$

$$\text{therefore,} \quad P (i/1) = \frac{\lambda_i}{\sum_{\text{All } i} \lambda_i} \quad (29)$$